

09-01-00

A

Please type a plus sign (+) inside this box → ☐

PTO/SB/05 (4/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

| | | |
|--|--|---|
| UTILITY PATENT APPLICATION TRANSMITTAL <small>(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))</small> | Attorney Docket No. | 1662-31400 (P00-3212) |
| | First Inventor or Application Identifier | Michael S. BERTONE et al. |
| | Title | Mechanism To Control The Allocation Of An N-Source ... Buffer |
| | Express Mail Label No. | EL705960437US |

| | |
|---|---|
| APPLICATION ELEMENTS <small>See MPEP chapter 600 concerning utility patent application contents.</small> | ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231 |
| 1. <input checked="" type="checkbox"/> * Fee Transmittal Form (e.g., PTO/SB/17) <small>(Submit an original and a duplicate for fee processing)</small> 2. <input checked="" type="checkbox"/> Specification [Total Pages 21] <small>(preferred arrangement set forth below)</small> - Descriptive title of the Invention (plus cover sheet) - Cross References to Related Applications - Statement Regarding Fed sponsored R & D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings (if filed) - Detailed Description - Claim(s) - Abstract of the Disclosure 3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) [Total Sheets 4] 4. Oath or Declaration [Total Pages 2] a. <input checked="" type="checkbox"/> Newly executed (original or copy) b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) <small>(for continuation/divisional with Box 16 completed)</small> i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b). | 5. <input type="checkbox"/> Microfiche Computer Program (Appendix) 6. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identity of above copies |
| ACCOMPANYING APPLICATION PARTS | |
| 7. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & document(s)) 8. <input type="checkbox"/> 37 C.F.R. § 3.73(b) Statement (when there is an assignee) <input type="checkbox"/> Power of Attorney 9. <input type="checkbox"/> English Translation Document (if applicable) 10. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations 11. <input type="checkbox"/> Preliminary Amendment 12. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized) 13. <input type="checkbox"/> * Small Entity Statement(s) <input type="checkbox"/> Statement filed in prior application, Status still proper and desired (PTO/SB/09-12) 14. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed) 15. <input type="checkbox"/> Other: | |
| NOTE FOR ITEMS 1 & 13: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28). | |

16. If a **CONTINUING APPLICATION**, check appropriate box, and supply the requisite information below and in a preliminary amendment

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No: _____

Prior application information. Examiner: _____ Group / Art Unit: _____

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Label 23505 or ☒ Correspondence address below
(Insert Customer No. or Attach bar code label here)

| | | | | | |
|---------|----------------------------|-----------|--------------|----------|--------------|
| Name | Jonathan M. Harris | | | | |
| | Conley, Rose & Tayon, P.C. | | | | |
| Address | PO Box 3267 | | | | |
| City | Houston | State | TX | Zip Code | 77253-3267 |
| Country | USA | Telephone | 713-238-8000 | Fax | 713-238-8008 |

| | | | |
|-------------------|---------------------------|-----------------------------------|-----------------|
| Name (Print/Type) | Jonathan M. Harris | Registration No. (Attorney/Agent) | 44,144 |
| Signature | <i>Jonathan M. Harris</i> | Date | August 31, 2000 |

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application Washington, DC 20231.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR UNITED STATES LETTERS PATENT

**MECHANISM TO CONTROL THE ALLOCATION
OF AN N-SOURCE SHARED BUFFER**

By:

Michael S. Bertone
119 Roundtop Road
Marlborough, MA 01752
Citizenship: U.S.A.

Richard E. Kessler
30 Thetland Dr.
Shrewsbury, MA 01545
Citizenship: U.S.A.

David H. Asher
415 Central Turnpike
Sutton, MA 01590
Citizenship: U.S.A.

Steve Lang
359 Taylor Road
Stow, MA 01775
Citizenship: U.S.A.

MECHANISM TO CONTROL THE ALLOCATION OF AN N-SOURCE SHARED BUFFER

5

CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to the following commonly assigned co-pending applications entitled:

“Scan Wheel – An Apparatus For Interfacing A High Speed Scan-Path With A Slow Speed Tester,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-23700; “Rotary Rule And Coherence Dependence Priority Rule,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27300; “Speculative Scalable Directory Based Cache Coherence Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27400;

“Scalable Efficient IO Port Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27500; “Efficient Translation Buffer Miss Processing For Applications Using Large Pages In Systems With A Large Range Of Page Sizes By Eliminating Page Table Level,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27600; “Fault Containment And Error Recovery Techniques In A Scalable Multiprocessor,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27700; “Speculative Directory Writes In A Directory Based CC-Non Uniform Memory Access Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27800; “Special Encoding Of Known Bad Data,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-27900; “Broadcast Invalidate Scheme,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-28000; “Mechanism To Keep All Pages Open In A DRAM Memory System,” Serial

No. _____, filed August 31, 2000, Attorney Docket No. 1662-28100; "Programmable
DRAM Address Mapping Mechanism," Serial No. _____, filed August 31, 2000, Attorney
Docket No. 1662-28200; "Mechanism To Enforce Memory Read/Write Fairness, Avoid Tristate
Bus Conflicts, And Maximize Memory Bandwidth," Serial No. _____, filed August 31,
5 2000, Attorney Docket No. 1662-29200; "An Efficient Address Interleaving With Simultaneous
Multiple Locality Options," Serial No. _____, filed August 31, 2000, Attorney Docket No.
1662-29300; Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-29400;
"Method And System For Absorbing Defects In High Performance Microprocessor With A Large
N-Way Set Associative Cache," Serial No. _____, filed August 31, 2000, Attorney Docket
10 No. 1662-29500; "A Method For Reducing Directory Writes And Latency In A High Performance,
Directory-Based, Coherency Protocol," Serial No. _____, filed August 31, 2000, Attorney
Docket No. 1662-29600; "Mechanism To Reorder Memory Read And Write Transactions For
Reduced Latency And Increased Bandwidth," Serial No. _____, filed August 31, 2000,
Attorney Docket No. 1662-30800; "Look-Ahead Mechanism To Minimize And Manage Bank
15 Conflicts In A Computer Memory System," Serial No. _____, filed August 31, 2000,
Attorney Docket No. 1662-30900; "Resource Allocation Scheme That Ensures Forward Progress,
Maximizes Utilization Of Available Buffers And Guarantees Minimum Request Rate," Serial
No. _____, filed August 31, 2000, Attorney Docket No. 1662-31000; "Input Data Recovery
Scheme," Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-31100; "Fast
20 Lane Prefetching," Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-
31200; "Mechanism For Synchronizing Multiple Skewed Source-Synchronous Data Channels
With Automatic Initialization Feature," Serial No. _____, filed August 31, 2000, Attorney
Docket No. 1662-31300; and "Chaining Directory Reads And Writes To Reduce DRAM

Bandwidth In A Directory Based CC-NUMA Protocol,” Serial No. _____, filed August 31, 2000, Attorney Docket No. 1662-31500, all of which are incorporated by reference herein.

**STATEMENT REGARDING FEDERALLY SPONSORED
RESEARCH OR DEVELOPMENT**

Not applicable.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to the distribution of buffer space between multiple sources. More particularly, the invention relates to a fair and efficient method of controlling allocation of a shared buffer pool.

Background of the Invention

In computer systems and networks, buffers are a convenient means of storing commands, requests, and data that are in transit from one location to another. Buffers may be used in a variety of applications. They may be used in network switching devices to temporarily hold data packets while network congestion subsides or while the switch determines the location to which the data must be forwarded. It is not uncommon for network switches to manage traffic for a plurality of sources. Buffers may also be used in memory and data allocation. An example of the latter would be a data read/write request buffer that must allocate requests from multiple sources. A common problem in systems using a shared buffer space is signal traffic that creates congestion and may lead to buffer overflow and monopolization by one or more of the buffer sources.

Ideally, a system comprising a buffer with multiple sources should accomplish several tasks. First, the system should not deliver data, commands, or requests to the buffer if the buffer

does not have any free space. This prevents data loss or packet drops which may require that the data packet be re-sent resulting in even greater bandwidth loss than simply holding the data until buffer space becomes available. Secondly, access to buffer space by the multiple sources should preferably be allocated in a fair manner. The sources should have fair access to the buffer so a source does not become backlogged while other sources are able to deliver data freely. This does not necessarily imply that the allocation needs to be equal for each source. For instance, one source may be given priority over the others. However, even in this scenario, it is important to prevent complete monopolization by the source that has priority.

One conventional solution to the problem of fair allocation of a shared buffer space is hard partitioning of the buffer space. For example, if a buffer has 16 data spaces and the buffer is shared among 4 sources, each source may be allocated four buffer slots. This method of allocation is certainly fair but may be horribly inefficient. If one of the sources has a string of data that ideally could be burst to the buffer, congestion may occur because the source only has four buffer spaces available. The remaining 12 buffers could be used to hold the burst of data, but instead may lie dormant because of the hard partitioning.

If prior knowledge exists about the type of traffic that can be expected from the sources, the hard partitioning may be altered to allocate more buffer space to one source or another. For instance, in the example given above, seven buffer spaces may be allocated to one source while the other three sources are allocated three spaces each. This allocation may alleviate congestion for the prioritized source, but does not prevent the burst congestion for of the other sources. In either case, hard partitioning tends to preclude use of at least a fixed percentage of the buffer space unless all sources are continuously accessing the buffer.

Another conventional solution to the problem of fairly and efficiently allocating buffer space is with stop and go commands issued by the buffer. In this type of system, the buffer is configured to keep track of available spaces within the buffer. During normal operation with light traffic, each source receives a "go" signal from the buffer indicating that buffer space is available.

5 As buffer space becomes limited, the buffer may send "stop" signals to the individual sources to halt data transmission to the buffer. This approach offers better use of the buffer space because the sources are not limited to a fixed percentage of the buffer space. However, some risk is involved in this type of embodiment because a finite time exists between the moment the buffer sends out a stop command and the moment the source receives and responds to the stop command. During
10 this finite time, it is possible for additional data to be transmitted to the buffer from the various sources. If the buffer was sufficiently close to being full and enough data was sent to the buffer before the stop commands were received by the sources, buffer overflow may occur and data may be lost. To prevent this, stop commands are usually sent well in advance of the buffer filling to capacity. Thus, if all buffer sources are bursting data to the buffer, the stop command is preferably
15 timed so that the sources stop sending data to the buffer before the buffer capacity is filled. Unfortunately, the side effect of sending the stop commands early is that the maximum capacity of the buffer will not be used when the buffer sources are not simultaneously sending bursts of data. The stop/go command solution to this buffer allocation problem is an improvement over the hard partitioning solution, but presents problems of either overflow or inefficient use of the whole
20 buffer.

It is desirable therefore to develop a fair and efficient means of allocating buffer space among several sources. The allocation scheme preferably prevents monopolization by any of the buffer sources. The allocation method also preferably takes advantage of the full capacity of the

buffer space without the possibility of buffer overflow. The system may advantageously be applied to a plurality of buffer sources and may also be applied to a variety of applications.

BRIEF SUMMARY OF THE INVENTION

5 The problems noted above are solved in large part by a method and apparatus for ensuring fair and efficient use of a shared memory buffer. The invention uses a credit-based allocation scheme to prevent monopolization by one or more sources and permits efficient use of the entire buffer. A preferred embodiment comprises a shared memory request buffer in a multi-processor computer system. The shared memory request buffer is used to store requests from different processors. Memory requests from a local processor are delivered to the local memory controller by a cache control unit. Requests from other processors are delivered to the memory controller by an interprocessor router. The memory controller allocates the memory requests in a shared buffer using a credit-based allocation scheme. The cache control unit and the interprocessor router are each assigned a number of credits. Each must pay a credit to the memory controller when a request is allocated to the shared buffer. If a source does not have any available credits, that source may not send a request to the shared buffer. The number of credits assigned to each source is sufficient to enable each source to deliver an uninterrupted burst of memory requests to the buffer without having to wait for credits to return from the buffer. The total number of credits assigned to the sources is preferably small compared to the overall size of the buffer. If the number of filled spaces in the shared buffer is below a threshold, the buffer immediately returns the credits to the source from which the credit and memory request arrived. If the number of filled spaces in the shared buffer is above a threshold, the buffer holds the credits and returns a single credit in a round-robin manner only when a space in the shared buffer becomes free. The buffer threshold is

the point when the number of free spaces available in the buffer is equal to the total number of credits assigned to the cache control unit and the interprocessor router. Since credits are not freely returned as the buffer gets full and since there are never any more credits available than spaces in the buffer, the buffer may reach capacity, but will not overflow.

5

BRIEF DESCRIPTION OF THE DRAWINGS

For a detailed description of the preferred embodiments of the invention, reference will now be made to the accompanying drawings in which:

Figure 1 shows a system diagram of a plurality of microprocessors coupled together;

Figures 2 shows a block diagram of a microprocessor memory controller in which the preferred embodiment is implemented;

Figure 3 shows a schematic representation of the credit-based allocation of the memory request buffer in the memory controller of Figure 2 operating in a light load condition; and

Figure 4 shows the memory request buffer of Figure 3 operating in a heavy load condition.

NOTATION AND NOMENCLATURE

Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, computer companies may refer to a component by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms “including” and “comprising” are used in an open-ended fashion, and thus should be interpreted to mean “including, but not limited to...”. Also, the term “couple” or “couples” is intended to mean either an indirect or direct electrical connection. Thus, if a first device couples to a second device,

that connection may be through a direct electrical connection, or through an indirect electrical connection via other devices and connections.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

5 The preferred embodiment of the shared buffer allocation scheme is directed toward application in a memory request buffer used in a multi-processor computer system. While the preferred embodiment may be implemented in a variety of applications, a detailed description in the context of the multi-processor computer system is provided herein. The descriptions herein are not intended to limit the scope of the claim limitations set forth in this specification.

10 Referring now to Figure 1, in accordance with the preferred embodiment of the invention, computer system 90 comprises one or more processors 100 coupled to a memory 102 and an input/output (“I/O”) controller 104. As shown, computer system 90 includes twelve processors 100, each processor coupled to a memory and an I/O controller. Each processor preferably includes four ports for connection to adjacent processors. The interprocessor ports are designated
15 “North,” “South,” “East,” and “West” in accordance with the well-known Manhattan grid architecture also known as a crossbar interconnection network architecture. As such, each processor 100 can be connected to four other processors. The processors on both ends of the system layout wrap around and connect to processors on the opposite side to implement a 2D torus-type connection. Although twelve processors 100 are shown in the exemplary embodiment
20 of Figure 1, any desired number of processors (*e.g.*, 256) can be included. For purposes of the following discussion, the processor in the upper, left-hand corner of Figure 1 will be discussed with the understanding that the other processors 100 are similarly configured in the preferred embodiment.

As noted, each processor preferably has an associated I/O controller 104. The I/O controller 104 provides an interface to various input/output devices such as disk drives 105 and 106, as shown in the lower, left-hand corner of Figure 1. Data from the I/O devices thus enters the 2D torus via the I/O controllers.

5 Each processor also, preferably, has an associated memory 102. In accordance with the preferred embodiment, the memory 102 preferably comprises RAMbus™ memory devices, but other types of memory devices can be used, if desired. The capacity of the memory devices 102 can be any suitable size. Further, memory devices 102 preferably are implemented as Rambus Interface Memory Modules (“RIMM”). To aid in the control of distributed memory in the
10 multiprocessor system, each processor includes a memory manager and directory structure for the local memory source. A preferred embodiment of the memory controller is shown in Figure 2.

In general, computer system 90 can be configured so that any processor 100 can access its own memory 102 and I/O devices, as well as the memory and I/O devices of all other processors in the system. Preferably, the computer system may have physical connections between each
15 processor resulting in low interprocessor communication times and improved memory and I/O device access reliability. If physical connections are not present between each pair of processors, a pass-through or bypass path is preferably implemented in each processor that permits accesses to a processor’s memory and I/O devices by another processor through one or more pass-through processors.

20 Referring now to Figure 2, each processor 100 preferably includes an L2 instruction and data cache control unit (“Cbox”) 170, an L2 data cache 180, a memory controller (“Zbox”) 190, and an interprocessor and I/O router unit (“Rbox”) 200. The following discussion describes each of these units.

Each of the functional units 170, 180, 190, 200 contains control logic that communicates with the control logic of other functional units as shown. Other functional units may exist within the processor 100, but have been omitted from Figure 1 for clarity in the description of the preferred embodiment. For example, the processor may further comprise integer and floating-point execution units, a memory reference unit, and an instruction fetch, issue and retire unit.

The L2 instruction and data cache control unit ("Cbox") 170 controls the L2 instruction and data cache 180 and handles data accesses from other functional units in the processor, other processors in the computer system, or any other devices that might need data out of the L2 cache.

The interprocessor and I/O router unit ("Rbox") 200 provides the interfaces to as many as four other processors and one I/O controller 104 (Figure 1). The inter-processor interfaces are designated as North ("N"), South ("S"), East ("E"), and West ("W") and provide two-way communication between adjacent processors.

Processor 100 preferably includes at least one RAMbus™ memory controller 190 ("Zbox"). The Zbox 190 controls 4 or 5 channels of information flow with the main memory 102 (Figure 1). The Zbox 190 preferably includes a back end 193, a middle mapper 192, and a front-end directory in flight table ("DIFT") 191. The middle mapper 192 maps the physical address into RAMbus™ device format by device, bank, row, and column. The middle mapper 192 also maintains an open-page table to track all open pages and to close pages on demand if bank conflicts arise. The mapper 192 also schedules RAMbus™ transactions such as timer-base request queues. The Zbox back end 193 preferably packetizes the address, control, and data into RAMbus™ format and provides the electrical interface to the RAMbus™ devices themselves.

The front-end DIFT 191 performs a number of functions such as managing the processor's directory-based memory coherency protocol, processing request commands from the Cbox 170 and

Rbox 200, sending forward commands to the Rbox 200, sending response commands to and receiving packets from the Cbox 170 and Rbox 200.

The DIFT 191 also comprises a shared request buffer for tracking up to thirty-two in-flight transactions. These transaction requests or packets are received from the Cbox 170 or the Rbox 200. When a request comes from either source, the DIFT must allocate an entry in the 32-space buffer. The requests from the Cbox 170 are from the local processor chip whereas requests from the Rbox 200 are off chip requests. Since each processor in the multi-processor system shown in Figure 1 is identical, a Cbox 170 request should carry the same priority as a request from the Rbox 200. It is important therefore to allocate space in the shared DIFT buffer in a fair and efficient manner.

Figure 3 represents a simplified schematic of the preferred embodiment of a credit allocation scheme for the DIFT buffer. The DIFT 191 preferably includes a shared request buffer 300 and credit counters 310, 320, 330. The Cbox 170 preferably includes a credit counter 340. Likewise, the Rbox includes a credit counter 350. The credit counters in the DIFT 191, Cbox 170 and Rbox 200 determine when the Cbox 170 and Rbox 200 may deliver a request to the DIFT 191. For both the Cbox 170 and Rbox 200, a request may only be delivered to the DIFT 191 when a credit is available in the local credit counter 340 and 350, respectively. A credit from Cbox counter 340 is spent (i.e., given up to the DIFT credit counter 320) by the Cbox 170 when a request is sent to the DIFT buffer 300. Similarly, a credit from Rbox counter 350 is spent (i.e., given up to the DIFT credit counter 330) by the Rbox 200 when a request is sent to the DIFT buffer 300.

In the preferred embodiment, credits are returned from the DIFT 191 to the Cbox 170 and Rbox 200 in two distinct manners. The method of returning credits depends on the region of operation for the DIFT buffer 300. The first occurs under light loads when the DIFT buffer 300 is

relatively empty. This condition is met when the number of occupied buffer spaces is below a threshold 360. In the example shown in Figure 3, the DIFT buffer 300 only has a few buffer spaces filled. In this situation, the DIFT 191 immediately returns credits to the source from which a request arrives. The Cbox 170 and Rbox 200 will therefore have all or nearly all of their credits in light load conditions.

The number of credits assigned to the Cbox 170 and Rbox 200 are based, in part, on credit round trip times. Since credits are returned immediately during light load situations, the total time required to transmit a credit from the Cbox 170 to the DIFT 191 and for the DIFT 191 to return the credit to the Cbox 170 may be determined. The number of credits given to the Cbox 170 is based not only on this round trip time, but also on the speed with which the Cbox 170 may deliver a burst of requests. The preferred embodiment gives enough credits so that if the Cbox 170 has enough memory requests from the processor, it may continuously deliver these requests without having to wait for credits to return from the DIFT 191. Consider for example that the Cbox 170 has a string of requests that need to be sent yet only has 4 credits. Ideally if the Cbox bursts these requests one after another as quickly as it can, then by the time the Cbox is ready to transmit the fifth request, the credit from the first request has preferably arrived back at the Cbox 170. This credit may then be used to transmit the fifth request. Subsequent requests may then be transmitted with other credits as they arrive back at the Cbox 170. This process may continue as long as the number of occupied spaces in the DIFT buffer remains under the threshold 360.

The number of credits given to the Rbox 200 is determined in the same way as for the Cbox 170. The round trip times between the two sources should not be assumed to be the same. A source with longer credit round trip times will necessarily require more credits to guarantee uninterrupted burst requests. Thus the number of credits assigned to the Rbox 200 will probably,

though not necessarily, differ from the number of credits given to the Cbox 170. In the system shown in Figure 3, the Cbox counter 340 holds four credits while the Rbox counter 350 holds five credits. The actual number of credits given to each source in the preferred embodiment is system specific and will depend on round trip times between source and buffer.

5 The second manner in which the DIFT 191 returns credits to the Cbox 170 and Rbox 200 occurs under heavier loads when the DIFT buffer 300 is relatively full. The threshold 360 between these two conditions is preferably defined as the difference between the size of the DIFT buffer 300 and the number of credits distributed among the sources. In the system shown in Figure 3, the threshold occurs when there are only nine buffer spaces left. When the DIFT buffer 300 is filled above this threshold, the DIFT 191 will hold credits until a buffer space is freed. This situation is depicted in Figure 4.

10 In Figure 4, the DIFT buffer 300 is filled above the threshold 360 and therefore, the DIFT 191 holds onto the credits spent by the Cbox 170 and Rbox 200. The credits held from the Cbox 170 and Rbox 200 are stored in counters 320 and 330, respectively. In this region of operation, the DIFT buffer counter 310 indicates the number of DIFT buffer spaces available. The threshold 360 is set so that the number of outstanding credits never exceeds the number of free spaces in the DIFT buffer 300. Thus, if the Cbox 170 and Rbox 200 send requests until all credits are gone, the DIFT buffer will be full. The preferred embodiment may therefore advantageously use the entire DIFT buffer 300 without running the risk of buffer overflow.

15 In this second region of operation, each time a space in the DIFT buffer 300 becomes available, credits are preferably returned to the Cbox 170 or Rbox 200. To ensure fairness, the credits are returned in a random, equally probable round-robin fashion to those sources which have spent credits. A suitable method of returning the credits may be to simply alternate credit returns.

Other embodiments exist where statistical methods are used to determine which source receives a credit. For instance, if priority needs to be given to the Cbox 170, the credit returns may be based on a random generator that is statistically skewed to return more credits to the Cbox 170 than the Rbox 200. The preferred embodiment however, returns credits on an equally likely basis to ensure fairness between the sources. It should also be noted that when the system operates in the first region of operation (during light loads), fairness between the sources is guaranteed because credits are returned immediately.

An alternative embodiment exists whereby the credit allocation scheme may be used to send request or response commands to the DIFT buffer 300. In this embodiment, the Cbox 170 and Rbox 200 are assigned credits as discussed above, but each must reserve a final credit for a response only. Thus, all credits in the Cbox 170 or Rbox 200 may be spent in issuing requests or responses to the DIFT buffer 300, but the last credit must be spent on a response.

The above discussion is meant to be illustrative of the principles and various embodiments of the present invention. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. For example, the shared buffer space may be extended to include three or more buffer sources. It is intended that the following claims be interpreted to embrace all such variations and modifications.

CLAIMS

What is claimed is:

- 1 1. A multi-processor computer system, comprising:
 - 2 a plurality of processors, each processor coupled to at least one memory cache, one cache
 - 3 control unit, and one interprocessor router;
 - 4 a memory coupled to each processor, each memory managed by a memory controller
 - 5 configured to accept memory requests from the plurality of processors; and
 - 6 at least one input/output device coupled to at least one processor;
 - 7 wherein the memory requests from a local processor are delivered to the memory controller
 - 8 by the cache control unit and wherein memory requests from other processors are delivered to the
 - 9 memory controller by the interprocessor router and wherein the memory controller allocates the
 - 10 memory requests in a shared buffer using a credit-based allocation scheme .
- 1 2. The computer system of claim 1 wherein:
 - 2 the cache control unit and the interprocessor router are each assigned a number of credits;
 - 3 at least one of said credits must be delivered by the cache control unit to the memory
 - 4 controller when a memory request is delivered by the cache control unit to the memory controller;
 - 5 and
 - 6 at least one of said credits must be delivered by the interprocessor router to the memory
 - 7 controller when a memory request is delivered by the interprocessor router to the memory
 - 8 controller;
 - 9 wherein if the number of filled spaces in the shared buffer is below a threshold, the buffer
 - 10 returns the credits to the source from which the credit and memory request arrived.

1 3. The computer system of claim 2 wherein:
2 wherein if the number of filled spaces in the shared buffer is above a threshold, the buffer
3 holds the credits and returns a credit in a round-robin manner to a source from which a credit has
4 been received only when a space in the shared buffer becomes free; and
5 wherein if a source has no available credits, that source cannot deliver a memory request to
6 the shared buffer.

4. The computer system of claim 2 wherein:
the number of credits assigned to the cache control unit and the interprocessor router is
sufficient to enable each source to deliver an uninterrupted burst of memory requests to the buffer
without having to wait for credits to return from the buffer.

5. The computer system of claim 4 wherein:
the number of credits available in the cache control unit and the interprocessor router are
stored and updated in counters located in the cache control unit and the interprocessor router; and
the number of credits spent by the cache control unit and the interprocessor router are
stored and updated in counters located in the shared buffer.

6. The computer system of claim 4 wherein:
the threshold is the point when the number of free spaces available in the buffer is equal to
the total number of credits assigned to the cache control unit and the interprocessor router.

1 7. A computer processor for use in a multi-processor system, comprising:
2 an associated memory;
3 a memory controller comprising a request buffer in a front-end directory in-flight table;
4 an L2 data cache;
5 an L2 instruction and data cache control unit configured to send request and response
6 commands from the processor to the memory controller;
7 at least one input/output device coupled to the processor; and
8 an interprocessor and I/O router unit configured to send request and response commands
9 from other processors to the memory controller;
10 wherein the L2 instruction and data cache control unit and interprocessor and I/O router
11 unit are assigned a number of credits and are configured to give up a credit to the directory in-flight
12 table each time a request or response command is sent to the request buffer and wherein if the
13 request buffer is filled below a buffer threshold, the directory in-flight table immediately returns
14 credits to the source from which the credit was received.

1 8. The computer processor of claim 7 wherein:
2 if the request buffer is filled above a buffer threshold, the directory in-flight table holds
3 credits and returns a credit to a source from which a credit was received only when a buffer space
4 is emptied; and
5 wherein if a source has no available credits, that source may not send a request or response
6 command to the request buffer and wherein if a source has one available credits, that source may
7 only send a response command to the request buffer.

1 9. The computer processor of claim 8 wherein:
2 the credits are returned to the sources which have given up credits to the directory in-flight
3 table in a random, equally probably manner.

1 10. The computer processor of claim 8 wherein:
2 the buffer threshold is the point above which the number of empty spaces in the request
3 buffer is equal to the total number of credits assigned to the L2 instruction and data cache control
4 unit and interprocessor and I/O router.

1 11. The computer processor of claim 8 wherein the directory in-flight table further comprises:
2 a counter to store and update the number of credits spent by the L2 instruction and data
3 cache control unit;
4 a counter to store and update the number of credits spent by the interprocessor and I/O
5 router; and
6 a counter to store and update the number of empty spaces in the request buffer when the
7 request buffer is filled above the buffer threshold;
8 wherein when the request buffer is filled above the buffer threshold, the directory in-flight
9 table holds credits and returns credits only when the number of empty spaces in the buffer
10 increases.

1 12. The computer processor of claim 8 wherein:
2 the number of credits available to the L2 instruction and data cache control unit and
3 interprocessor and I/O router is stored and updated by counters in each unit.

1 13. The computer processor of claim 8 wherein:

2 the number of credits available to the L2 instruction and data cache control unit and
3 interprocessor and I/O router is determined by the round trip time required to send a credit to and
4 receive a credit from the directory in-flight table;

5 wherein the number of credits given to each source is sufficient to allow each source to
6 send an uninterrupted sequence of request or response commands to the directory in-flight table
7 without delays caused by waiting for credits to return from the directory in-flight table.

8 14. A method of allocating space in a shared buffer, comprising:

9 assigning credits to each source that sends data packets to the shared buffer; and

10 requiring each source to spend a credit each time that source sends a data packet to the
11 shared buffer;

12 wherein if the number of empty buffer spaces is larger than a buffer threshold, immediately
13 paying the credit back to the source from which the credit and data were sent; and

14 wherein if the number of empty buffer spaces is smaller than the buffer threshold, holding
15 the credit until a buffer space becomes empty and then paying a credit back to a source from which
16 a credit was sent.

1 15. The method of claim 14, wherein:

2 when the number of empty buffer spaces is smaller than the buffer threshold and a buffer
3 space becomes empty, returning a credit in a random, equally probably manner to one of the
4 sources which have spent credits held by the buffer.

1 16. The method of claim 14, wherein:
2 when the number of empty buffer spaces is smaller than the buffer threshold and a buffer
3 space becomes empty, returning a credit in a random, statistically skewed manner to one of the
4 sources which have spent credits held by the buffer.

1 17. The method of claim 14, further comprising:
2 assigning a minimum number of credits to each source that is sufficient to allow each
3 source to send a continuous sequence of data packets without waiting for returned credits.


1 18. The method of claim 14, further comprising:
2 preventing a source from delivering a data packet to the shared buffer if that source has no
3 available credits.


1 19. The method of claim 14, further comprising:
2 setting the buffer threshold equal to the number of total credits assigned to all the sources.

1 20. The method of claim 14, further comprising:
2 using a counter in each source and a counter for each source coupled to the buffer to track
3 spent and paid back credits.

ABSTRACT

A method and apparatus for ensuring fair and efficient use of a shared memory buffer. A preferred embodiment comprises a shared memory buffer in a multi-processor computer system. Memory requests from a local processor are delivered to a local memory controller by a cache control unit and memory requests from other processors are delivered to the memory controller by an interprocessor router. The memory controller allocates the memory requests in a shared buffer using a credit-based allocation scheme. The cache control unit and the interprocessor router are each assigned a number of credits. Each must pay a credit to the memory controller when a request is allocated to the shared buffer. If the number of filled spaces in the shared buffer is below a threshold, the buffer immediately returns the credits to the source from which the credit and memory request arrived. If the number of filled spaces in the shared buffer is above a threshold, the buffer holds the credits and returns the credits in a round-robin manner only when a space in the shared buffer becomes free. The number of credits assigned to each source is sufficient to enable each source to deliver an uninterrupted burst of memory requests to the buffer without having to wait for credits to return from the buffer. The threshold is the point when the number of free spaces available in the buffer is equal to the total number of credits assigned to the cache control unit and the interprocessor router.





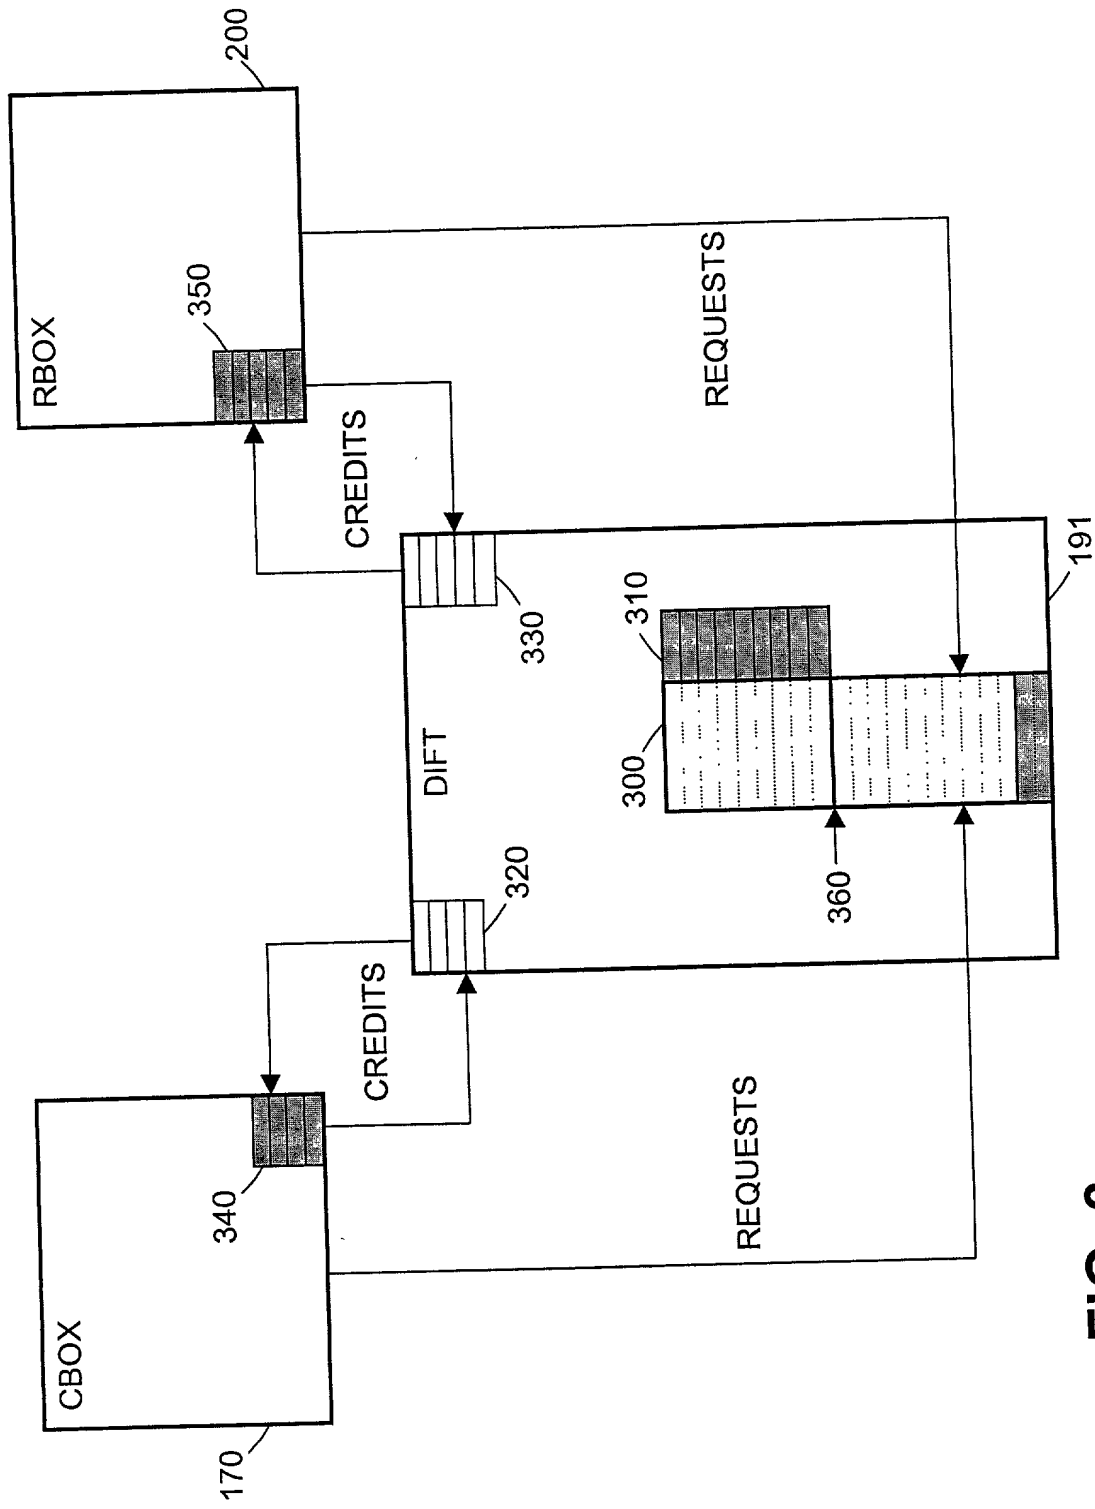


FIG. 3

FIG. 4 is a block diagram of a system for processing requests and credits. The system includes a CBOX (170) and an RBOX (200) connected to a central DIFT (191). The CBOX (170) contains a CREDIT buffer (340) and a REQUEST buffer (360). The RBOX (200) contains a CREDIT buffer (350) and a REQUEST buffer (360). The DIFT (191) contains a CREDIT buffer (320) and a REQUEST buffer (330). Arrows indicate the flow of credits and requests between the CBOX, RBOX, and DIFT.

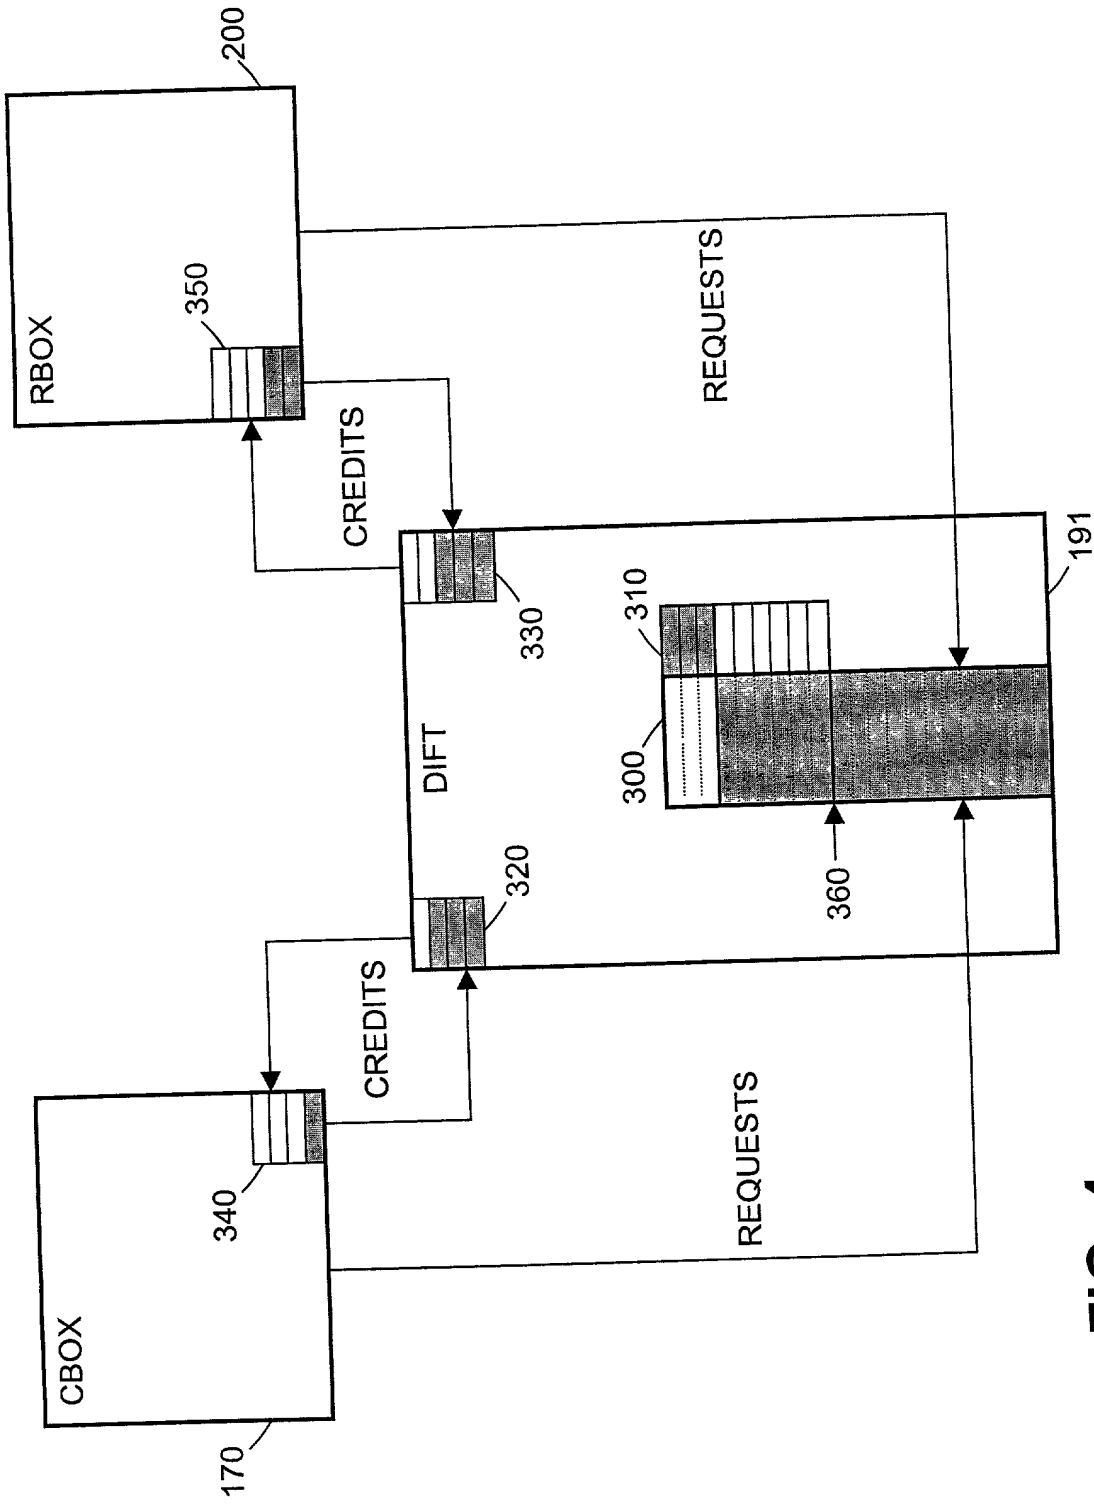


FIG. 4

DECLARATION

SOLE/JOINT INVENTOR
ORIGINAL/SUBSTITUTE/CIP

As a below named inventor, I hereby declare that: my residence, post office address, and citizenship are as stated below next to my name. I believe I am the original, first, and sole inventor (if only one name is listed below) or a joint inventor (if plural inventors are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: "MECHANISM TO CONTROL THE ALLOCATION OF AN N-SOURCE SHARED BUFFER", as described in the specification attached.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above; that I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof, or patented or described in any printed publication in any country before my or our invention thereof or more than one year prior to this application; that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representative or assigns more than twelve months prior to this application; and that I acknowledge the duty to disclose information of which I am aware which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations § 1.56(a). Such information is material when it is not cumulative to information already of record or being made of record in the application, and

- (1) it establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) it refutes, or is inconsistent with, a position the applicant has taken or may take in:

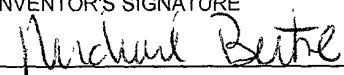

- (i) opposing an argument of unpatentability relied on by the Office, or
- (ii) asserting an argument of patentability.

I hereby claim foreign priority benefits under Title 35, United States Code § 119 of any foreign application(s) for patent or inventor's certificates listed below and have also identified below any foreign application(s) having a filing date before that of the application(s) on which priority is claimed:

| COUNTRY | APPLICATION NUMBER | DATE OF FILING | PRIORITY CLAIMED UNDER 35 USC 119 |
|---------|--------------------|----------------|--|
| | | | <input type="checkbox"/> YES <input type="checkbox"/> NO |

I hereby claim the benefit under Title 35 United States Code § 120 of any United States application(s) listed below and, insofar as any subject matter of any claim of this application is not disclosed in the prior United States Application, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations § 1.56(a) which occurred between the filing date of the prior application and the national PCT international filing date of this application:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

| | | |
|--|---|-------------|
| FULL NAME OF SOLE OR FIRST INVENTOR | INVENTOR'S SIGNATURE | DATE |
| Michael S. BERTONE |  | 8/28/00 |
| RESIDENCE | | CITIZENSHIP |
| 119 Roundtop Road, Marlborough, Massachusetts, 01752 | | U.S.A. |
| POST OFFICE ADDRESS | | |
| SAME AS ABOVE | | |
| FULL NAME OF SECOND JOINT INVENTOR | INVENTOR'S SIGNATURE | DATE |
| Richard E. KESSLER |  | 8/28/00 |
| RESIDENCE | | CITIZENSHIP |
| 30 Thetland Drive, Shrewsbury, Massachusetts, 01545 | | U.S.A. |
| POST OFFICE ADDRESS | | |
| SAME AS ABOVE | | |

Express Mail Label No. EL705960437
Attorney Docket No. 1662-31400
Client Docket No. P00-3212

DECLARATION

SOLE/JOINT INVENTOR
ORIGINAL/SUBSTITUTE/CIP

| | | |
|--|--------------------------|-------------|
| FULL NAME OF THIRD JOINT INVENTOR | INVENTOR'S SIGNATURE | DATE |
| David H. ASHER | <i>David H. Asher</i> | 8/28/00 |
| RESIDENCE | | CITIZENSHIP |
| 415 Central Turnpike, Sutton, Massachusetts, 01590 | | U.S.A. |
| POST OFFICE ADDRESS | | |
| SAME AS ABOVE | | |
| FULL NAME OF FOURTH JOINT INVENTOR | INVENTOR'S SIGNATURE | DATE |
| Steve LANG | <i>Steve Elbert Lang</i> | 8/28/00 |
| RESIDENCE | | CITIZENSHIP |
| 359 Taylor Road, Stow, Massachusetts, 01775 | | U.S.A. |
| POST OFFICE ADDRESS | | |
| SAME AS ABOVE | | |